

Deep reinforcement learning and simulation as a path toward precision medicine

CASIS 2018

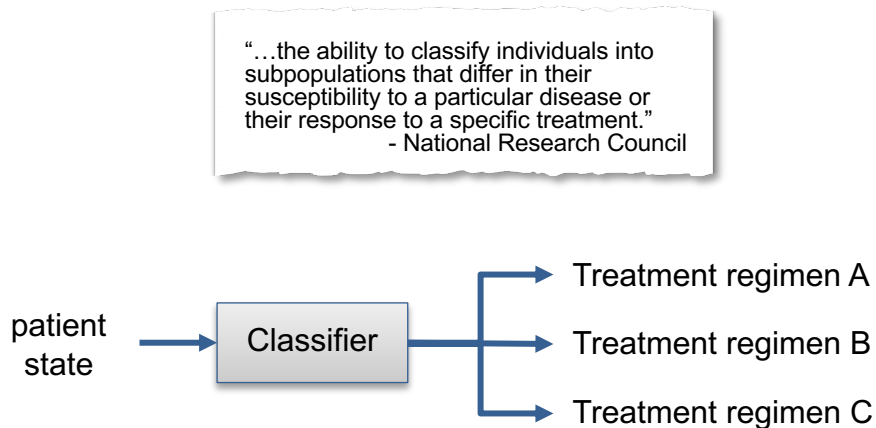
Brenden Petersen, Jiachen Yang, Will Grathwohl,
Chase Cockrell, Claudio Santiago, Thomas Desautels,
Gary An, Dan Faissol



Precision medicine as a control problem

Traditional precision medicine


Classify then treat



- Viewed as a classification task
- Therapies are static and non-adaptive

Proposed vision


Dynamic, feedback control





- Viewed as an optimal control task
- Therapies are dynamic and adaptive
 - Dependent upon patient trajectory

The need for simulation

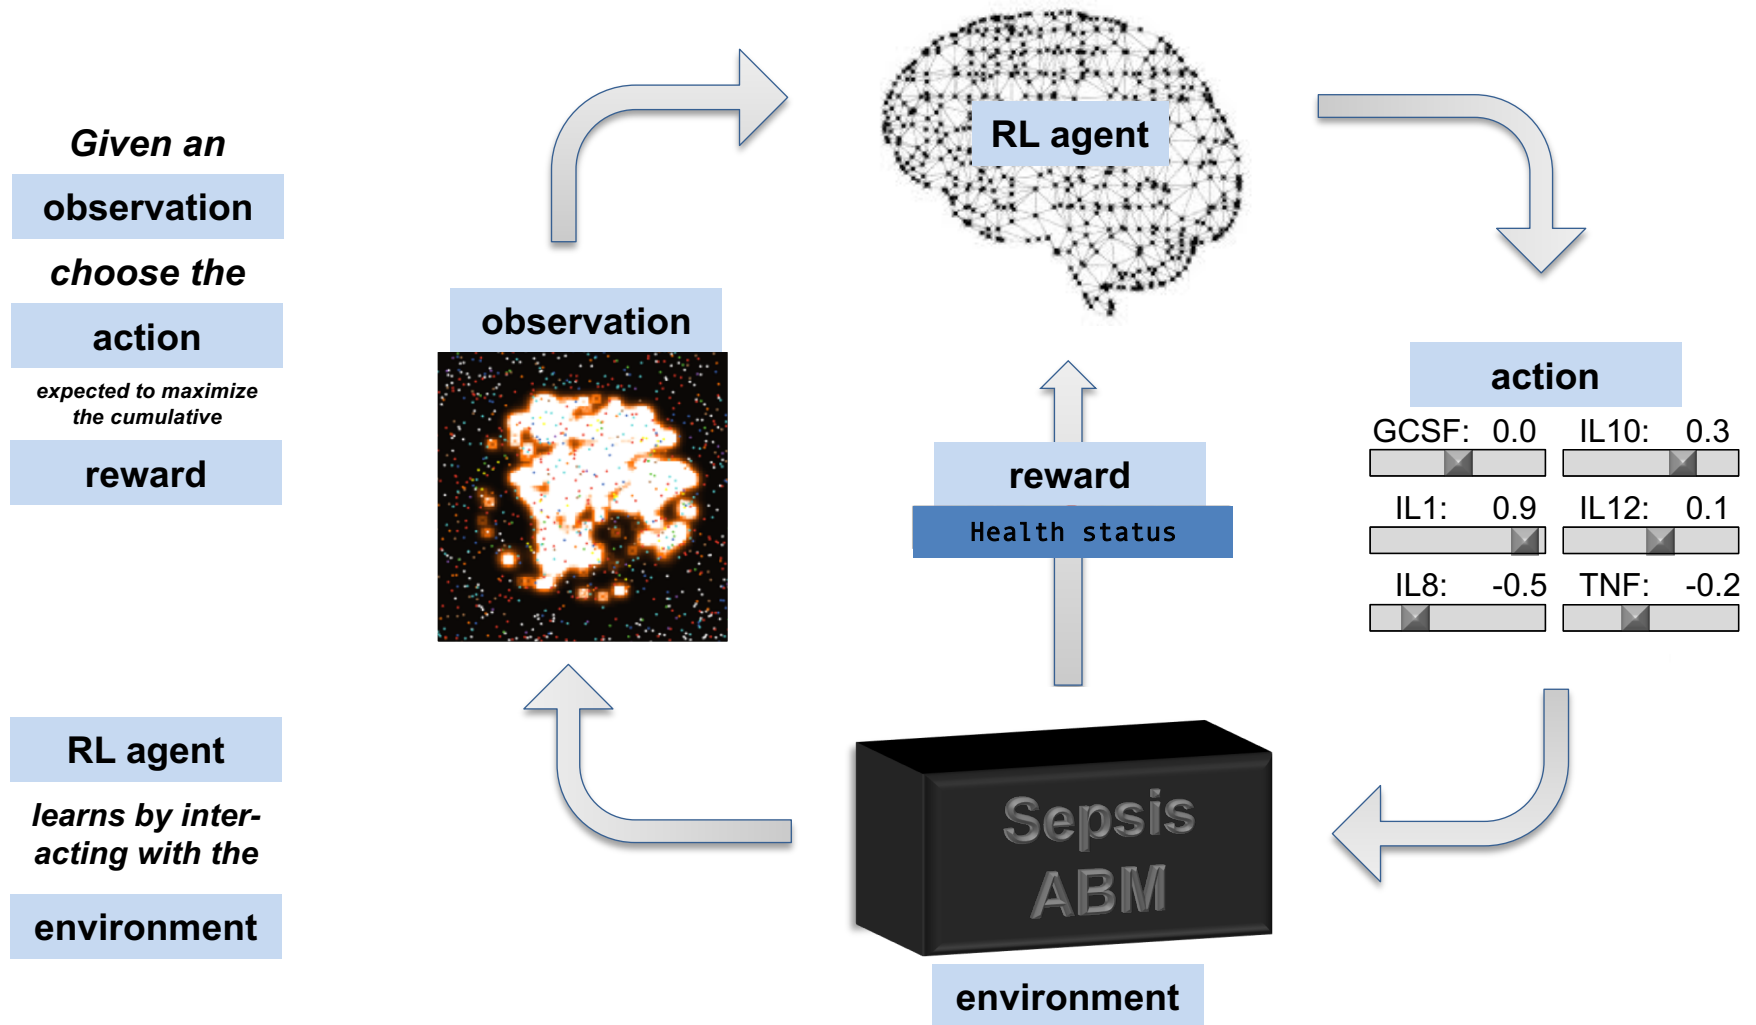
- Many control approaches use existing data to retrospectively learn control policies
- Simulation enables virtual experimentation: going beyond what has been tried
- Recent advances in optimal control have enabled learning controllers for complex, high-dimensional simulations

	Learning controllers using...	
	Clinical Data 	Biological Simulation
<i>Scope</i> of interventions	Limited to what's already been tried	Able to explore new interventions and/or combinations
<i>Interpretability</i> of interventions	Limited by statistical power of existing data	Limited only by computation
<i>Dimensionality</i> of interventions	Low-dimensional, discrete (e.g. 1 – 2 drugs, 3 doses)	High-dimensional, continuous
<i>Dynamics</i> of interventions	Typically static	Dynamic, adaptive

Sepsis agent-based simulation – Demo



Reinforcement learning (RL)



Problem Formulation: Observation Space

Observation Space

large, spatial

Cytokine level + cell counts
at each grid point

Size: $\mathbb{R}^{101 \times 101 \times 21 \times N}$


Clinically unrealistic with
today's technology

small, aggregate

Aggregate cytokine levels +
cell counts (non-spatial)

Size: $\mathbb{R}^{21 \times N}$

Clinically plausible from
blood tests



Problem Formulation: Action Space

GCSF: 0.0



IL1: 0.9



IL8: -0.5



IL10: 0.3



IL12: 0.1



TNF: -0.2



⋮

Action Space

large, continuous

Differentially control all cytokines at once

Size: $[-1, 1]^{14}$

Clinically plausible with multi-channel infusion pump

small, discrete

Augment or inhibit by a fixed amount;
One cytokine at a time

Size: 29

Clinically plausible

Problem Formulation: Reward Signal

- The simulation naturally provides only sparse, binary rewards: life/death

$$r_{\text{outcome}} = \lambda_+[\text{heal}] - \lambda_-[\text{die}]$$

- To aid learning, we added two terms to the reward signal
 1. Potential-based reward shaping term
 - Helps guide the RL agent toward “good” states without altering the optimal policy

$$r_\phi = \lambda_\phi (\text{damage}(s) - \text{damage}(s'))$$


2. A penalty for taking actions
 - Regularizer; promotes conservative actions

$$r_a = -\lambda_a \|a\|_1$$

- Final reward signal: $r(s, a, s') = r_{\text{outcome}} + r_\phi + r_a$

Unique challenges of the sepsis environment


Failed to solve using human experience, genetic algorithms, and classify → control approaches



Challenge	Atari 2600	MuJoCo	Sepsis
High-dimensional state	✓	✓	✓
High-dimensional actions		✓	✓
Sparse rewards	sometimes		✓
Long time horizons			✓
Computationally expensive			✓
Unsolvable by humans			✓
Stochastic	None	None	High
Each episode has different dynamics			✓

Training the DRL agent


- Environment is “solved” by 2,500 episodes
- Distinct “phases” of learning



Evaluating the learned policy


- Mortality rate under learned policy

- Trained patient: 46% \rightarrow 0%
- Across 500 patients: 49% \rightarrow 0.8%




- Clinical insight

- IL-1 (pro-inflammatory) is unregulated early and suppressed late
- Suppression comes later for patients with larger initial infections

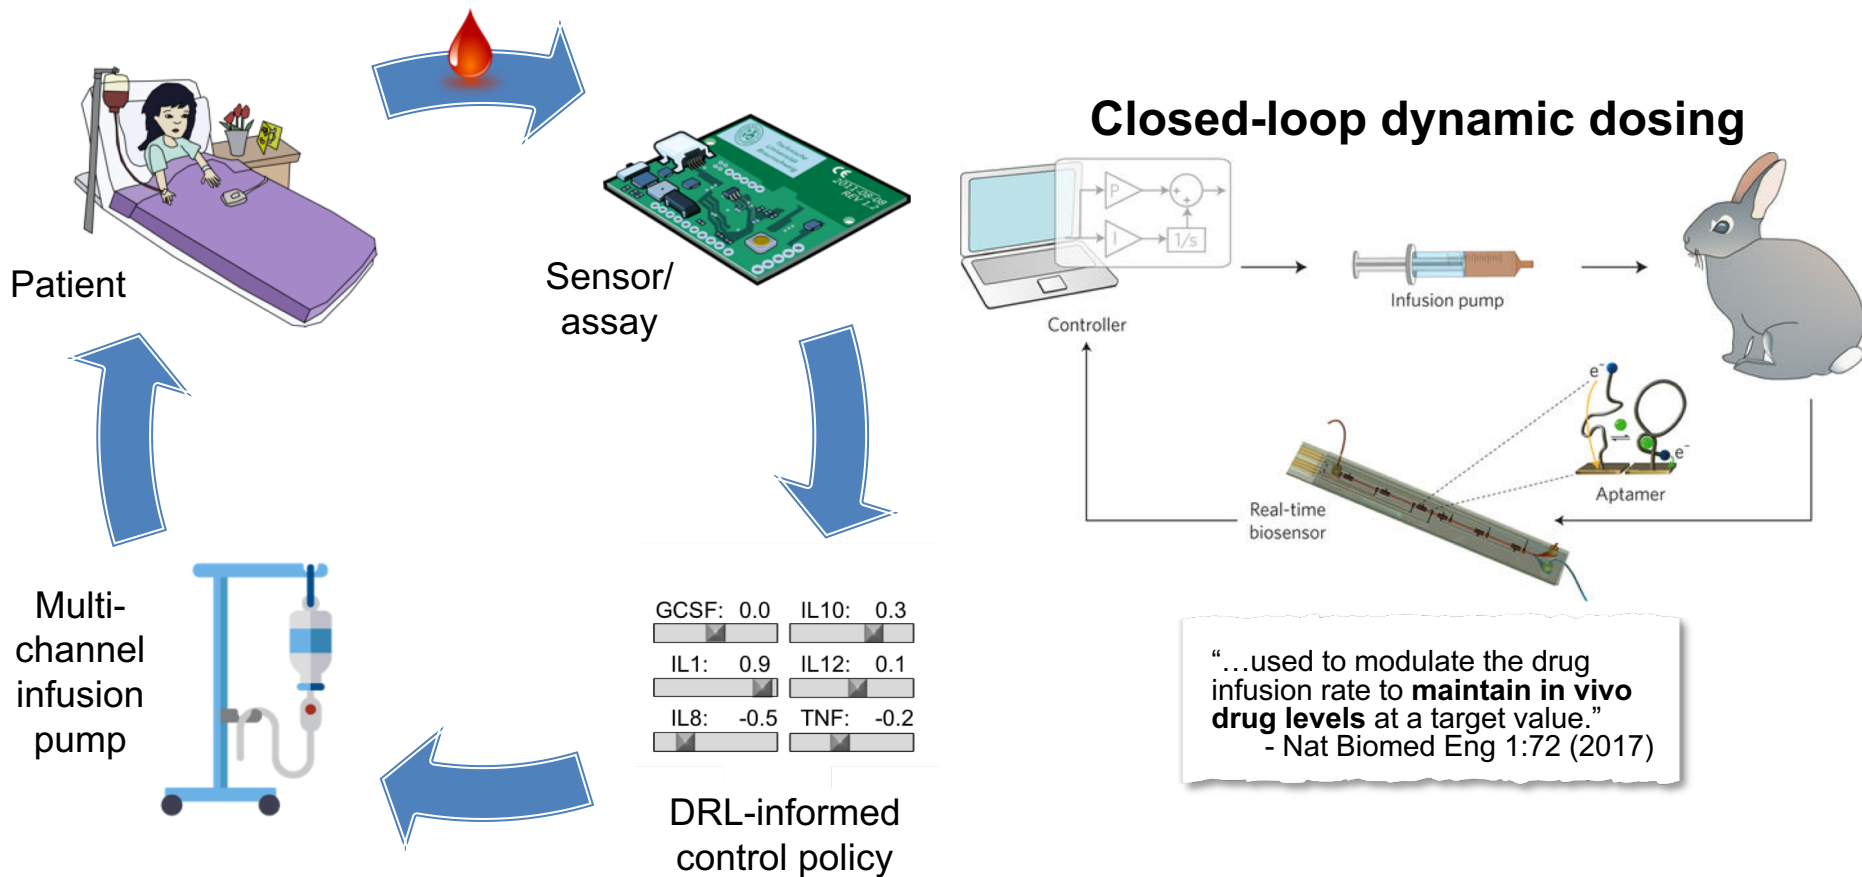


Next steps: Improving clinical plausibility

- Tradeoff between *controllability* and *clinical relevance*



Long-term vision: Closed-loop control system



<https://openclipart.org/>
<https://www.mediware.com/home-care/blog/new-legislation-help-home-infusion-patients/>

Thank you!

See Tom Desautel's poster!

